arXiv:1503.05951v1 [cs.LG] 19 Mar 2015

# Rank Subspace Learning for Compact Hash Codes

**Kai Li** KAILI@EECS.UCF.EDU
University of Central Florida, 4000 Central Florida Blvd., Orlando, FL 32826 USA

**Guojun Qi** GUOJUN.QI@UCF.EDU
University of Central Florida, 4000 Central Florida Blvd., Orlando, FL 32826 USA

**Jun Ye** JYE@EECS.UCF.EDU
University of Central Florida, 4000 Central Florida Blvd., Orlando, FL 32826 USA

**Kien A. Hua** KIENHUA@EECS.UCF.EDU
University of Central Florida, 4000 Central Florida Blvd., Orlando, FL 32826 USA

## Abstract

The era of Big Data has spawned unprecedented interests in developing hashing algorithms for efficient storage and fast nearest neighbor search. Most existing work learn hash functions that are numeric quantizations of feature values in projected feature space. In this work, we propose a novel hash learning framework that encodes feature's rank orders instead of numeric values in a number of optimal low-dimensional ranking subspaces. We formulate the ranking subspace learning problem as the optimization of a piecewise linear convex-concave function and present two versions of our algorithm: one with independent optimization of each hash bit and the other exploiting a sequential learning framework. Our work is a generalization of the Winner-Take-All (WTA) hash family and naturally enjoys all the numeric stability benefits of rank correlation measures while being optimized to achieve high precision at very short code length. We compare with several state-of-the-art hashing algorithms in both supervised and unsupervised domain, showing superior performance in a number of data sets.

## 1. Introduction

Massive amount of social multimedia data are being generated by billions of users every day. The advent of multimedia big data presents a number of challenges and opportunities for research and development of efficient storage, indexing and retrieval techniques. Hashing is recognized by many researchers as a promising solution to the above Big Data problem, thus attracting significant amount of research in the past few years (Datar et al., 2004; Kulis & Darrell, 2009; Tschopp & Diggavi, 2009; Weiss et al., 2009). Most hashing algorithms encode high-dimensional data into binary codes by quantizing numeric projections (Norouzi & Fleet, 2011; Liu et al., 2011; 2014; 2012). In contrast, hashing schemes based on feature's ranking order (i.e. comparisons) are relatively underresearched and will be the focus of this paper.

Ranking-based hashing, such as Winner-Take-All (WTA) (Yagnik et al., 2011) and Min-wise Hashing (MinHash) (Broder et al., 2000), ranks the random permutation of input features and uses the index of maximal/minimal feature dimensions to encode a compact representation of the input features. The benefit of ranking-based hashing lies in the fact that these algorithms are insensitive to the magnitude of features, and thus are more robust against many types of random noises universal in real applications ranging from information retrieval (Salakhutdinov & Hinton, 2007), image classification (Fan, 2013) to object recognition (Torralba et al., 2008). In addition, the magnitude-independence also makes the resultant hash codes scale-invariant, which is critical to compare and align the features from heterogeneous spaces, e.g., revealing the multi-modal correlations (Li et al., 2014).

Unfortunately, the existing ranking-based hashing is data-agnostic. In other words, the obtained hash codes are not learned by exploring the intrinsic structure of data distribution, making it suboptimal in its efficiency of coding the input features with compact codes of minimal length. For

example, WTA encodes the data with the indices of the maximum dimensions chosen from a number of random permutations of input features. Although WTA has generated leading performances in many tasks (Dean et al., 2013; Yagnik et al., 2011; Li et al., 2014), it is constrained in the sense that it only ranks the existing features of input data, while incapable of combining multiple features to generate new feature subspaces to rank. A direct consequence of such limitation is that this sort of ranking-based hashing usually needs a very large number of permutations and rankings to generate useful codes, especially with a high dimensional input feature space (Yagnik et al., 2011).

To address this challenge, we abandon the use of ranking random permutations of existing features in ranking-based hashing algorithms. Instead, we propose to generate compact ranking-based hashing codes by learning a set of new subspaces and ranking the newly projected features in these subspaces. At each step, an input data is encoded by the index of the maximal value over the projected points onto these subspaces. The subspace projections are jointly optimized to generate the ranking indices that are most discriminative to the metric structure and/or the data labels. Then a vector of codes are iteratively generated to represent the input data from the maximal indices over a sequence of sets of subspaces.

This method generalizes ranking-based hashing from restricted random permutations to perform encoding by ranking a set of arbitrary subspaces learned by mixing multiple original features. This greatly extends its flexibility so that much shorter bits can be generated to encode input data, while retaining the benefits of noise insensitivity and scale invariance inherent in such algorithms.

In the remainder of this paper, we first review the related hashing algorithms in Section 2, then the rank subspace hash learning problem is formulated and solved in Section 3. In Section 4, we present an improved learning algorithm based on sequential learning. The experimental results are presented in Section 5 and the paper is concluded in Section 6.

## 2. Related Work

Since the focus of this paper is on data-dependent hashing, we limit our review to two categories of this line of research – unsupervised and supervised hashing. Interested reader can refer to (Bondugula, 2013) and (Wang et al., 2014) for a comprehensive review of this research area.

The most representative works on unsupervised hashing includes Spectral Hashing (Weiss et al., 2009) and kernelized variant Locality Sensitive Hashing (Kulis & Grauman, 2009). In detail, SH learns linear projections through an eigenvalue decomposition so that distance between pairs of similar training samples are minimized in the projected subspaces when these samples are binarized. Similarly, KLSH also makes use of an eigensystem solution, but it manipulates data items in the kernel space in an effort to generalize LSH to accommodate arbitrary kernel functions. Binary Reconstructive Embedding (BRE) (Kulis & Darrell, 2009) explicitly minimizes the reconstruction error between the input space and Hamming space to preserve the metric structure in input space, which has demonstrated improved performance over SH and LSH. Iterative Quantization (ITQ) (Gong & Lazebnik, 2011) iteratively learns uncorrelated hash bits to minimize quantization error between hash bodes and the dimension-reduced data. (Wang et al., 2012) presents a sequential projection learning method that fits the eigenvector solution into a boosting framework and uses pseudo labels in learning each hash bit. More recently, Anchor Graph Hashing (AGH) (Liu et al., 2011) and Discrete Graph Hashing (DGH) (Liu et al., 2014) use anchor graphs to capture the neighborhood structure inherent in a given dataset and adopt a discrete optimization procedure to achieve nearly balanced and uncorrelated hash bits.

On the other hand, supervised hashing methods take advantage of the data labels to learn data-dependent hash functions. It has been shown that supervised hashing methods are good at incorporating data labels to learn more discriminative hash codes in many tasks. For example, (Salakhutdinov & Hinton, 2007) uses Restricted Boltzman Machine (RBM) to learn nonlinear binary hash codes for document retrieval and demonstrates better precision and recall than other methods. Similar deep learning hash methods have also been applied to the task of image retrieval in very large databases (Torralba et al., 2008). However, deep learning methods typically need large data sets, cost long training times and have been outperformed by other methods exclusively designed for learning hashing functions. For instance, (Norouzi & Fleet, 2011) proposes the Minimal Loss Hashing (MLH) with a structural SVM-like formulation and minimizes the loss-adjusted upper bound of a hinge-like loss function defined on pairwise similarity labels. The resulting hash codes have shown to give superior performance over the state-of-the-art. This method is further extend to minimize loss functions defined with triplet similarity comparisons (Norouzi et al., 2012). Similarly, (Li et al., 2013) also learns hash functions based on triplet similarity. On the contrary, the formulation is a convex optimization within the large-margin learning framework rather than structural SVM.

Recently, (Fan, 2013) theoretically proved the convergence properties of arbitrary sequential learning algorithms and proposed the Jensen Shannon Divergence (JSD) sequential learning method with a multi-class classification formulation. Supervised Hashing with Kernels (KSH) is another sequential learning algorithm. This method maps

the data to compact hash codes by minimizing Hamming distances of similar pairs and maximizing that of dissimilar pairs simultaneously (Liu et al., 2012). The sequential part of our algorithm is similar to Boosting Similarity Sensitive Coding (BSSC) (Shakhnarovich et al., 2003) and Forgiving Hash (FH) (Baluja & Covell, 2008), both of which treat each hash bit as a week classifier and learn a series of hash functions in a AdaBoost framework. However, the rank-based hash function we learn at each step is significantly different from that in BSSC and FH, resulting in completely different objective functions and learning steps. Existing hashing schemes based on rank orders (e.g. (Tschopp & Diggavi, 2009), (Pele & Werman, 2008) and (Ozuysal et al., 2007)) are mostly restricted to approximating nearest neighbors given a distance metric to speed up large scale lookup. To the best of our knowledge, there has been no previous work explicitly exploiting the rank-based hash functions in a supervised hash learning setting.

# 3. Formulation

## 3.1. Winner-Take-All Hashing

The WTA hashing is a subfamily of hashing functions introduced by (Yagnik et al., 2011). WTA is specified by two parameters: the number of random permutations $L$ and the window size $K$. Each permutation $\boldsymbol{\pi}$ rearranges the entries of an input vector $\mathbf{x} \in \mathbb{R}^d$ to $\mathbf{x}_\pi$ in the order specified by $\boldsymbol{\pi}$. Then the index of the maximum dimension of the feature among the first $K$ elements of $\mathbf{x}_\pi$ is used as the hash code. This process is repeated $L$ times, resulting in a $K$-nary hash code of length $L$, which can be compactly represented using $L \times \lceil \log_2 K \rceil$ bits.

WTA is considered as a ranking-based hashing algorithm, which uses the rank order among permuted entries of a vector rather than their values of features. This property has given WTA certain degree of stability to perturbations in numeric values. Thus the WTA hash codes usually generate more robust metric structure to measure the similarity between input vectors than other types of hash codes which often contain inherent noises from quantizing the input feature spaces. With theoretical soundness, however, the hash codes generated by WTA often must be sufficiently long to represent the original data in high fidelity.

This is caused by twofold limitations: (1) the entries of input vectors are permuted in a random fashion before the comparison is applied to find the largest entry out of the first $K$ ones; (2) the comparison and the ranking are restricted to be made between the original features. The random permutations are very inefficient to find the most discriminative entries to compare the similarity between the input vectors, and the restriction of the ranking to original features is too strong to generate the compact representations. In the next, we relax the two limitations.

## 3.2. Rank Subspace Hashing

Rather than randomly permuting the input data vector $\mathbf{x}$, we project it onto a set of $K$ one-dimensional subspaces. Then the input vector is encoded by the index of the subspace that generates the largest projected value. In other words, we have

$$h(\mathbf{x}; \mathbf{W}) = \arg \max_{1 \le k \le K} \mathbf{w}_k^T \mathbf{x}, \tag{1}$$

where $\mathbf{w}_k \in \mathbf{R}^d, 1 \le k \le K$ are vectors specifying the subspace projections, and $\mathbf{W} = [\mathbf{w}_1, \mathbf{w}_2, \cdots, \mathbf{w}_K]^T$.

We use a linear projection to map an input vector into subspaces to form its hash code. At first glance, this idea is similar to the family of learning-based hashing algorithms based on linear projection (Datar et al., 2004; Norouzi & Fleet, 2011). However, different from these existing algorithms, the proposed method instead ranks the obtained subspaces to encode each input vector with the index of the dimension with the maximum value. This makes the obtained hash codes highly nonlinear to the input vector, invariant to the scaling of the vector, as well as insensitive to the input noises to a larger degree than the linear hashing codes. In this paper, we name this method Rank Subspace Hashing (RSH) to distinguish it from the other compared methods.

WTA is a special case of the RSH algorithm, if we restrict the projections onto $K$ axis-aligned linear subspaces, i.e., $\mathbf{w}_k$ is set to a column vector $\mathbf{e}_k$ randomly chosen from an identity matrix $\mathbf{I}$ of size $d \times d$.

RSH extends WTA by relaxing the axis aligned linear subspaces in (1) to arbitrary $K$-dimensional linear subspaces in $\mathbb{R}^d$. Such relaxation greatly increases the flexibility to learn a set of subspaces to optimize the hash codes resulting from the projections to these subspaces.

Now our objective boils down to learn hash functions characterized by the projections $\mathbf{W}$ as in Eq. (1). Specifically, let $\mathcal{D}$ be the set of $N$ $d$-dimensional data points $\{\mathbf{x}_i\}_{i=1}^N$ and let $\mathcal{S} = \{s_{ij}\}_{1 \le i,j \le N}$ be the set of pair-wise similarity labels satisfying $s_{ij} \in \{0, 1\}$, where $s_{ij} = 1$ means the pair $(\mathbf{x}_i, \mathbf{x}_j)$ is similar and vice verse. The pair-wise similarity labels $\mathcal{S}$ can be obtained either from the nearest neighbors in a metric space or by human annotation that denotes whether a pair of data points come from the same class.

Given a similarity label $s_{ij}$ for each training pair, we can define an error incurred by a hash function like (1) below

$$e(h_i, h_j, s_{ij}) = \begin{cases} \rho I(h_i \neq h_j), & s_{ij} = 1 \\ \lambda(1 - I(h_i \neq h_j)), & s_{ij} = 0 \end{cases} \tag{2}$$

where $I(\cdot)$ is the indicator function outputting 1 when the condition holds and 0 otherwise, $h_{i(j)}$ is $h(\mathbf{x}_{i(j)}; \mathbf{W})$ for short, and $\rho$ and $\lambda$ are two hyper-parameters that penalize false negative and false positive respectively.

The learning objective is to find $\mathbf{W}$ to minimize the cumulative error function over the training set:

$$E(\mathbf{W}) = \sum_{1 \le i \le j \le N} e(h_i, h_j, s_{ij}) \tag{3}$$

Note that $\mathbf{W}$ factors into the above objective function because both $h_i$ and $h_j$ are a function of $\mathbf{W}$.

### 3.3. Reformulation

The above objective function is straightforward to formulate, but hard to optimize because it involves the indicator function and $\arg\max$ function which are typically non-convex and highly discontinuous. Motivated by (Norouzi & Fleet, 2011), we reformulate the objective function and seek a piecewise linear upper bound of $E(\mathbf{W})$.

First, the hash function in (1) can be equivalently reformulated as

$$\begin{aligned} &\mathbf{h}(\mathbf{x}; \mathbf{W}) = \arg\max_{\mathbf{g}} \mathbf{g}^T \mathbf{W} \mathbf{x}, \\ &\text{subject to } \mathbf{g} \in \{0,1\}^K, \mathbf{1}^T \mathbf{g} = 1, \end{aligned} \tag{4}$$

which outputs an 1-of-$K$ binary code $\mathbf{h}$ for an input feature vector $\mathbf{x}$. The constraint enforces there must exist and only exist a nonzero entry of 1 in the resultant hash code. We enforce this constraint in the following optimization problems without meaning it explicitly to avoid notational clutter. It is easy to find the equivalence to the hashing function (1): the only nonzero element in $\mathbf{h}$ encodes the index of dimension with the maximum value in $\mathbf{W}\mathbf{x}$.

Given a pairwise similarity label $s_{ij}$ between two vectors $\mathbf{x}_i$ and $\mathbf{x}_j$, $\mathbf{h}_i$ and $\mathbf{h}_j$ are their hash codes obtained by solving the $\arg\max$ problem (i.e. $\mathbf{h}(\mathbf{x}_i; \mathbf{W})$ and $\mathbf{h}(\mathbf{x}_j; \mathbf{W})$). Then the error function (1) can be upper bounded by

$$\begin{aligned} e(\mathbf{h}_i, \mathbf{h}_j, s_{ij}) \le &\max_{\mathbf{g}_i, \mathbf{g}_j} [e(\mathbf{g}_i, \mathbf{g}_j, s_{ij}) + \mathbf{g}_i^T \mathbf{W} \mathbf{x}_i + \mathbf{g}_j^T \mathbf{W} \mathbf{x}_j] \\ &- \mathbf{h}_i^T \mathbf{W} \mathbf{x}_i - \mathbf{h}_j^T \mathbf{W} \mathbf{x}_j \end{aligned}$$

This inequality is easy to prove by noting that the following inequality

$$\begin{aligned} &\max_{\mathbf{g}_i, \mathbf{g}_j} [e(\mathbf{g}_i, \mathbf{g}_j, s_{ij}) + \mathbf{g}_i^T \mathbf{W} \mathbf{x}_i + \mathbf{g}_j^T \mathbf{W} \mathbf{x}_j] \\ &\ge e(\mathbf{h}_i, \mathbf{h}_j, s_{ij}) + \mathbf{h}_i^T \mathbf{W} \mathbf{x}_i + \mathbf{h}_j^T \mathbf{W} \mathbf{x}_j \end{aligned}$$

With the above upper bound of error function, we seek to solve the MinMax problem of minimizing the following function with respect to $\mathbf{W}$

$$\begin{aligned} \Omega(\mathbf{W}) = \sum_{1 \le i < j \le N} \{ &\max_{\mathbf{g}_i, \mathbf{g}_j} [e(\mathbf{g}_i, \mathbf{g}_j, s_{ij}) + \mathbf{g}_i^T \mathbf{W} \mathbf{x}_i + \mathbf{g}_j^T \mathbf{W} \mathbf{x}_j] \\ &- \mathbf{h}_i^T \mathbf{W} \mathbf{x}_i - \mathbf{h}_j^T \mathbf{W} \mathbf{x}_j \} \end{aligned}$$

### 3.4. Optimization

Consider $\mathbf{W}$ is fixed. The first step is a discrete optimization problem that is guaranteed to have global optimal solution. Specifically, given the values of $\mathbf{W}\mathbf{x}_{i(j)}$, the RSH codes $\mathbf{h}_{i(j)}$ in the second and third term of (5) can be found straightforwardly in $O(K)$. For the adjusted error $e(\mathbf{g}_i, \mathbf{g}_j, s_{ij}) + \mathbf{g}_i^T \mathbf{W} \mathbf{x}_i + \mathbf{g}_j^T \mathbf{W} \mathbf{x}_j$ of the first term in the square bracket, it is not hard to derive its maximum value can be obtained by scanning the elements in matrix $[m_{kl}]_{K \times K}$, defined as

$$m_{kl} = \begin{cases} y_i^{(k)} + y_j^{(l)} + \lambda(1 - s_{ij}) & \text{if } k = l \\ y_i^{(k)} + y_j^{(l)} + \rho s_{ij} & \text{otherwise} \end{cases} \tag{5}$$

where $y_i^{(k)}$ is the $k^{th}$ element of $\mathbf{W}\mathbf{x}_i$. Assuming the $(k^*, l^*)$th element of the above matrix achieves the maximum value, the maxima $(\mathbf{g}_i^*, \mathbf{g}_j^*)$ of the adjusted error are 1-of-$K$ binary vectors with the $k^*$th and the $l^*$th dimension set to 1. The above procedure can be computed in $O(K^2)$. Since $K$ is normally very small (e.g. 2 to 8), the above discrete optimization problem can be computed efficiently.

Now consider the optimization of $\mathbf{W}$. Fixing the maxima $(\mathbf{g}_i^*, \mathbf{g}_j^*)$ of the first term, and the RSH codes $\mathbf{h}_i$ and $\mathbf{h}_j$ in (4), $\mathbf{W}$ can be updated in the direction of the negative gradient

$$-\frac{\partial \Omega(\mathbf{W})}{\mathbf{W}} = \sum_{i,j} (\mathbf{h}_i - \mathbf{g}_i^*) \mathbf{x}_i^T + (\mathbf{h}_j - \mathbf{g}_j^*) \mathbf{x}_j^T \tag{6}$$

Batch update can be made using (6) when the training data can be loaded into the memory all at once. Otherwise, $\mathbf{W}$ can also be done in an online fashion with one training pair at a time, leading to the following iterative learning procedure

$$\mathbf{W} \leftarrow \mathbf{W} + \eta \big[ (\mathbf{h}_i - \mathbf{g}_i^*) \mathbf{x}_i^T + (\mathbf{h}_j - \mathbf{g}_j^*) \mathbf{x}_j^T \big], \tag{7}$$

where $\eta$ is the learning rate.

The learning algorithm is shown as Algorithm 1. The algorithm learns $L$ projection matrices by starting with different random initializations from Gaussian distribution. Because the convex-concavity nature of the objective function, the solutions have multiple local minima. This is a desired property in our application, because each local minimum, corresponding to a RSH function, reflects a distinct perspective of ranked subspaces underlying the training examples. In addition, each hash function is learned independently and thus can be done in parallel. The convergence

**Algorithm 1** Rank Subspace Learning

**Input:** data $[\mathbf{x}_i]$, pairwise similarity labels $[s_{ij}]$, length of hash code $L$, subspace dimension $K$
**for** $l = 1$ **to** $L$ **do**
  Initialize $\mathbf{w}_k, 1 \leq k \leq K$ from Gaussian distribution
  **repeat**
    Pick a pair $(\mathbf{x}_i, \mathbf{x}_j)$ and compute $\mathbf{h}_i$, $\mathbf{h}_j$, $\mathbf{g}_i^*$, $\mathbf{g}_j^*$
    Update projection matrix $\mathbf{W}$ according to
      $\mathbf{W} \leftarrow \mathbf{W} + \eta[(\mathbf{h}_i - \mathbf{g}_i^*)\mathbf{x}_i^T + (\mathbf{h}_j - \mathbf{g}_j^*)\mathbf{x}_j^T]$
  **until** Convergence
**end for**

of the learning algorithm has been explored and empirically studied in (McAllester et al., 2010; Norouzi & Fleet, 2011).

## 4. The Sequential Learning

In Algorithm 1, since each hash function is learned independently, the entire hash code may be suboptimal. This is because different random starting points may lead to the same local minima, resulting in redundant hash bits. In order to maximize the information contained in a $L$-bit hash code, we propose to learn the hash functions sequentially so each hash function can provide complementary information to previous ones.

In order to motivate our sequential learning algorithm, we can view each hash bit as a week classifier that assigns similarity labels to an input pair, and the obtained ensemble classifier is related with the Hamming distance between hashing codes. Formally, each week classifier corresponding to the $l^{th}$ bit is

$$sim_l(\mathbf{x}_i, \mathbf{x}_j) = 1 - H_m(\mathbf{h}(\mathbf{x}_i; \mathbf{W}_l), \mathbf{h}(\mathbf{x}_j; \mathbf{W}_l)) \quad (8)$$

Where $H_m(x, y) = I(x \neq y)$ is the bitwise Hamming distance, and $\mathbf{W}_l$ is the projection matrix for this bit. Then, the Hamming distance between two $L$-bit hash codes can be seen as the vote of an ensemble of $L$ week classifiers on them. Clearly, the sequential learning problem naturally fits into the AdaBoost framework.

The AdaBoost-based sequential learning algorithm is shown in Algorithm 2. In detail, a sampling weight $\alpha_{ij}^{(l)}$ is assigned to each training pair and is updated before training each new hash function. In particular, pairs that are misclassified by the current hash function will be given more weight in training the next hash function. The projection matrix is updated in the similar online fashion as in (7) but weighted by the sampling weight.

When all the hash functions have been trained, the voting results of the related week classifiers are fused with a

**Algorithm 2** Sequential Rank Subspace Learning

**Input:** data $[\mathbf{x}_i]$, pairwise similarity labels$[s_{ij}]$, length of hash code $L$, subspace dimension $K$
**Initialize:** set all the sampling weights $\{\alpha_{ij}\}$ to 1
**for** $l = 1$ **to** $L$ **do**
  Initialize $\mathbf{W}_l$ from Gaussian distribution
  **repeat**
    Pick a pair $(\mathbf{x}_i, \mathbf{x}_j)$ and compute $\mathbf{h}_i$, $\mathbf{h}_j$, $\mathbf{g}_i^*$, $\mathbf{g}_j^*$ based on the current estimate of $\mathbf{W}_l$;
    Update projection matrix $\mathbf{W}_l$ according to
    $\mathbf{W}_l \leftarrow \mathbf{W}_l + \eta\alpha_{ij}^{(l)}\left[(\mathbf{h}_i - \mathbf{g}_i^*)\mathbf{x}_i^T + (\mathbf{h}_j - \mathbf{g}_j^*)\mathbf{x}_j^T\right]$;
  **until** Convergence
  Compute the weighted errors

$$\epsilon_l = \frac{\sum_{i,j} \alpha_{ij}^{(l)} e(\mathbf{h}_i, \mathbf{h}_j, s_{ij})}{\sum_{i,j} \alpha_{ij}^{(l)}}$$

  Evaluate the quantity

$$\theta_l = \ln\left\{\frac{1 - \epsilon_l}{\epsilon_l}\right\}$$

  Update the pair weighting coefficients using

$$\alpha_{ij}^{(l+1)} \propto \alpha_{ij}^{(l)} \exp\{\theta_l e(\mathbf{h}_i, \mathbf{h}_j, s_{ij})\}$$

  Normalize the sampling weights such that $\sum_{i,j} \alpha_{ij}^{(l+1)} = \sum_{ij} \alpha_{ij}^{(l)}$.
**end for**

weighted combination

$$sim(\mathbf{x}_i, \mathbf{x}_j) = \sum_{l=1}^{L} \theta_l (1 - H_m(\mathbf{h}(\mathbf{x}_i; \mathbf{W}_l), \mathbf{h}(\mathbf{x}_j; \mathbf{W}_l)), \quad (9)$$

where $\theta_l$ are the weighted training error of the $l^{th}$ hash function.

We name this AdaBoost-inspired sequential learning by Sequential RSH (SRSH), in contrast to the RSH algorithm with independently composed hash codes.

## 5. Experiments

### 5.1. Dataset and Compared Methods

In order to evaluate the proposed hashing approaches, Rank Subspace Hashing (RSH) and Sequential Rank Space Hashing (SRSH), we use three well-known datasets: LabelMe and Peekaboom, two collections of images represented as 512D Gist vectors designed for object recognition tasks; and MNIST, a corpus of handwritten digits in $24 \times 24$ greyscale image. The above datasets are assembled by (Kulis & Darrell, 2009) and also used in (Norouzi & Fleet, 2011).

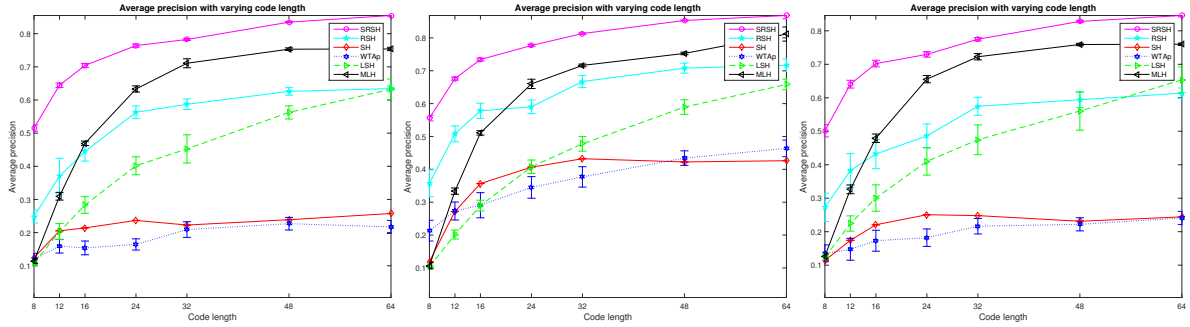


(a) LabelMe (b) MNIST (c) Peekaboom

*Figure 1.* Average precision with varying hash code length.

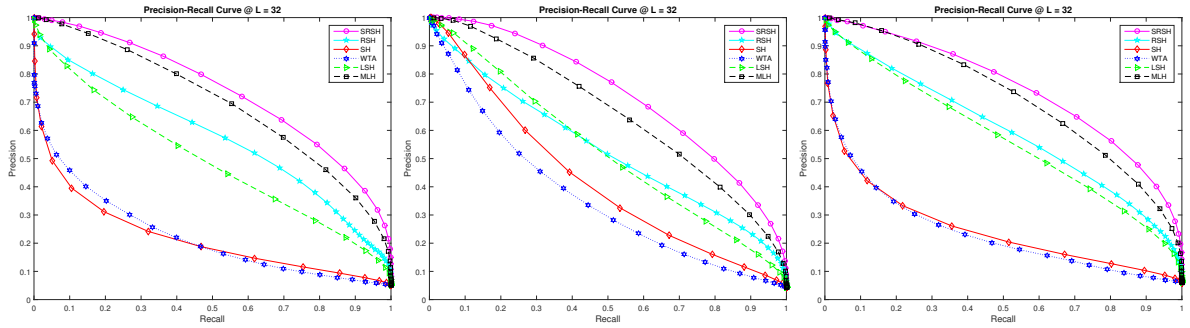


(a) LabelMe (b) MNIST (c) Peekaboom

*Figure 2.* Precision-recall curve when code length L = 32.

Following the settings of (Norouzi & Fleet, 2011), we randomly picked 1000 points for training and a separate set of 3000 points as test queries. The groundtruth neighbors for test queries are defined by thresholding the Euclidean distance such that each query point has an average of 50 neighbors. Similarly, we define the neighbors and non-neighbors of each data point in the training set in order to create the similarity matrix. All the datasets are mean-centered and normalized prior to training and testing. Some methods (e.g. SH) often perform better after dimensionality reduction, we therefore apply PCA to all datasets and retain the top 40 directions for a fair comparison.

For comparison, we choose several state-of-the-art methods: Minimal Loss Hashing (MLH (Norouzi & Fleet, 2011)), Spectral Hashing (SH (Weiss et al., 2009)), Locality Sensitive Hashing (LSH (Datar et al., 2004)), and Winner-Take-All (WTA (Yagnik et al., 2011)) hash. For MLH and SH, we use the publicly available source code provided by their original authors, while we implemented our own version of LSH and WTA since they are rather straightforward to implement. Those methods cover both supervised (e.g., MLH) and unsupervised (e.g., SH) hashing as well as data-agnostic ones (e.g., LSH and WTA), and are considered most representative in their own category.

### 5.2. Methodology

In evaluating Approximate Nearest Neighbor (ANN) search, two methods are frequently adopted in the literature, that is, hash table based lookup and Hamming distance based kNN search. We use both methods in our evaluation. In hash table lookup, the hash code is used to index all the points in a database, and the data points with the same hash key fall into the same bucket. Typically, hash buckets that fall within a Hamming ball of radius $R$ (i.e. the hash code differs by only 2 or 3 bits) of the target query are considered to contain relevant query results. A big advantage of hash table lookup lies in that it can be done in constant time. In contrast, Hamming distance based kNN search performs a standard kNN searching procedure based on Hamming distance which involves a linear scan of the entire database. However, since Hamming distance can be computed efficiently, the kNN search in Hamming space is also very fast in practice.

In our experiments, we evaluate the retrieval quality by setting $R = 2$ and 3 in the hash table lookup and $k = 50$ and 100 in Hamming distance based kNN search. For

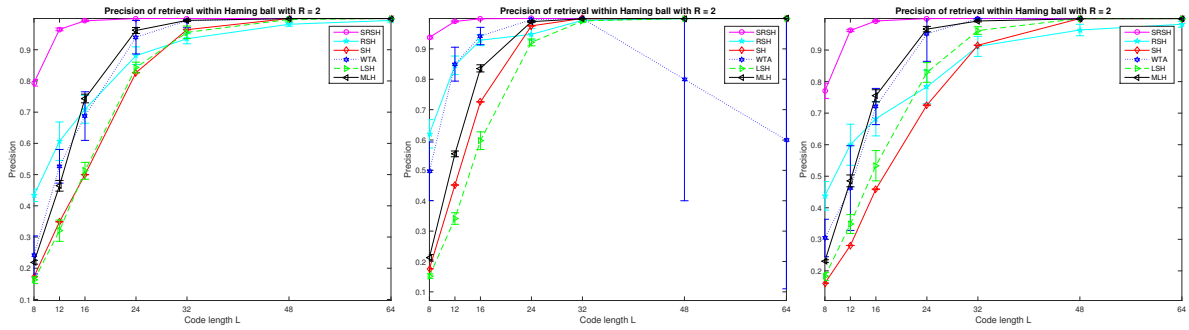


(a) LabelMe (b) MNIST (c) Peekaboom

*Figure 3.* Precision of retrieval within a Hamming ball of radius R = 2.

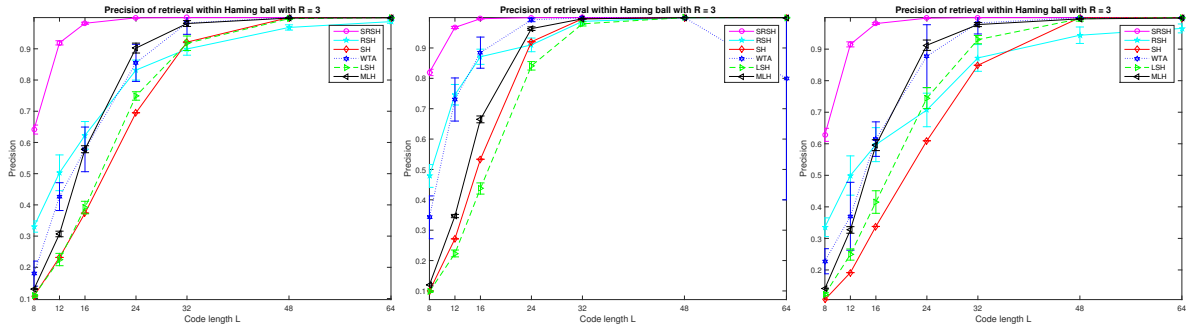


(a) LabelMe (b) MNIST (c) Peekaboom

*Figure 4.* Precision of retrieval within a Hamming ball of radius R = 3.

both evaluation protocols, we compute the retrieval precision that is defined as the percentage of true neighbors among those returned by the query. The precision reflects the quality of hash codes to a large extent and it can be critical for many applications. In addition, we also evaluate the average precision for hash table lookup by varying $R$, which approximates the area under precision-recall curve. For all benchmarks (unless otherwise specified), we run every algorithm 10 independent times and report the mean and the standard deviation.

As for parameter settings, MLH requires a loss scaling factor $\epsilon$, and two loss function hyper-parameters $\rho$ and $\lambda$. We follow the practice of MLH and perform cross-validation on a number of combinations to get the best performance at each code length. Similarly, our algorithm also has three hyper-parameter, that is, the subspace dimension $K$ and two error term hyper-parameters as defined in (2). A similar cross validation procedure is used to find the best model. For WTA, we use the polynomial kernel extension and set window size $K = 4$ and polynomial degree $p = 4$, as suggested by (Yagnik et al., 2011). SH and LSH are essentially parameter free and therefore do not require special handling.

### 5.3. Results

Figure 1 shows the average precision using different hash code length. We aim to compare the performance of different hashing methods in generating compact hash code, therefore the code length is restricted below 64. It can be observed from the Figure 1 that the average precision of almost all methods increases monotonically when codes become longer, which is reasonable since longer codes retains more information of original data. The only exception to this trend is SH whose performance doesn't increase or even slightly drops after exceeding certain number of bits (e.g. 24 to 32). This can be explained by fact that unsupervised learning methods tend to overfit more easily with longer codes, which is consistent with the observation by (Wang et al., 2012).

We note that RSH shows significant improvement over WTA, another representative ranking-based hashing algorithm, as a result of the generalization of projection directions and the supervised learning process. Compared with RSH, SRSH further boosts the performance with large gains across all the tested datasets, demonstrating the effectiveness of the sequential learning method. In general, SRSH achieves the best performance, with about $10\%$

lead over MLH. We also note that both of our algorithms demonstrate exceptional performance with extremely short code (e.g. of length less than 12) as a result of using rank order encoding.

In addition to the average precision, we also show a more detailed precision-recall profile when the code length $L$ is fixed to 32 in Figure 2. In the precision-recall curve, better performance is shown by larger area under the curve. Again, both of our algorithms perform significantly better than WTA with SRSH consistently being the best, which is consistent with the previous results.

The results of hash table lookups are shown in Fig. 3 and Fig. 4, for $R = 2$ and $R = 3$ respectively. As explained in the previous section, precision alone is more critical than average precision that is an overall evaluation of both precision and recall. Therefore, the results in Figure 3 and Figure 4 can be more important for such applications. In those tests, rank order based techniques (i.e. WTA, RSH and SRSH) generally perform better than numeric value based hashing schemes because of certain degree of resilience to numeric noises/perturbations. For example, although both WTA and LSH are based on data-agnostic random methods, WTA clearly outperforms LSH for most of the tests, which is similar to the results obtained in ((Yagnik et al., 2011)). However, we find that WTA sometimes fails to retrieve any neighbor within a small Hamming ball, resulting in large standard deviation in precision at large code length (e.g. Fig. 3(b) and 4(b)). This is a natural result of applying randomness to a highly selective hash function. Such limitation is effectively addressed by providing certain supervision in obtaining the hash functions. Therefore, both RSH and SRSH produce more stable results than WTA, as demonstrated by the consistently smaller standard deviations. Overall, SRSH performs the best in all the tests, again demonstrating its superiority in generating high quality hash codes.

The last group of experiments is the Hamming distance based kNN search, where we evaluate the precision of true neighbors among the 50 and 100 nearest neighbors measured by Hamming distance. As shown in Figure 5 and Figure 6, the results are similar to those in the hash table lookups, except that there are no missed retrievals for any of the compared algorithms because all queries are guaranteed to return the specified number of results. The proposed algorithms both give competitive results as compared with the others.

## 6. Conclusion

In this paper, a new reformulation of the Winner-Take-All hashing scheme is first presented. Based on this formulation, we propose a novel hash learning objective that aims to optimize a number of low-dimensional linear subspaces for high quality rank order-based hash encoding. A simple yet effective learning algorithm is then provided to optimize the objective function, leading to a number of optimal rank subspaces. The effectiveness of the proposed learning method in addressing the limitations of WTA is verified in a number of experiments. We also embed our learning method into a sequential learning framework that pushes the performance of the basic learning algorithm even further. Extensive experiments on several well-known datasets demonstrated our superior performance over state-of-the-art.

## References


Baluja, Shumeet and Covell, Michele. Learning to hash: forgiving hash functions and applications. *Data Mining and Knowledge Discovery*, 17(3):402–430, 2008.

Bondugula, Sravanthi. Survey of hashing techniques for compact bit representations of images. *Ph.D. dissertation,*, abs/1408.2927, 2013.

Broder, Andrei Z, Charikar, Moses, Frieze, Alan M, and Mitzenmacher, Michael. Min-wise independent permutations. *Journal of Computer and System Sciences*, 60 (3):630 – 659, 2000. ISSN 0022-0000.

Datar, Mayur, Immorlica, Nicole, Indyk, Piotr, and Mirrokni, Vahab S. Locality-sensitive hashing scheme based on p-stable distributions. In *Proceedings of the twentieth annual symposium on Computational geometry*, pp. 253–262. ACM, 2004.

Dean, Thomas, Ruzon, Mark A, Segal, Mark, Shlens, Jonathon, Vijayanarasimhan, Sudheendra, and Yagnik, Jay. Fast, accurate detection of 100,000 object classes on a single machine. In *Computer Vision and Pattern Recognition (CVPR), 2013 IEEE Conference on*, pp. 1814–1821. IEEE, 2013.

Fan, Lixin. Supervised binary hash code learning with jensen shannon divergence. In *Computer Vision (ICCV), 2013 IEEE International Conference on*, pp. 2616–2623. IEEE, 2013.

Gong, Yunchao and Lazebnik, Svetlana. Iterative quantization: A procrustean approach to learning binary codes. In *Computer Vision and Pattern Recognition (CVPR), 2011 IEEE Conference on*, pp. 817–824. IEEE, 2011.

Kulis, Brian and Darrell, Trevor. Learning to hash with binary reconstructive embeddings. In *Advances in neural information processing systems*, pp. 1042–1050, 2009.

Kulis, Brian and Grauman, Kristen. Kernelized locality-sensitive hashing for scalable image search. In *Computer*

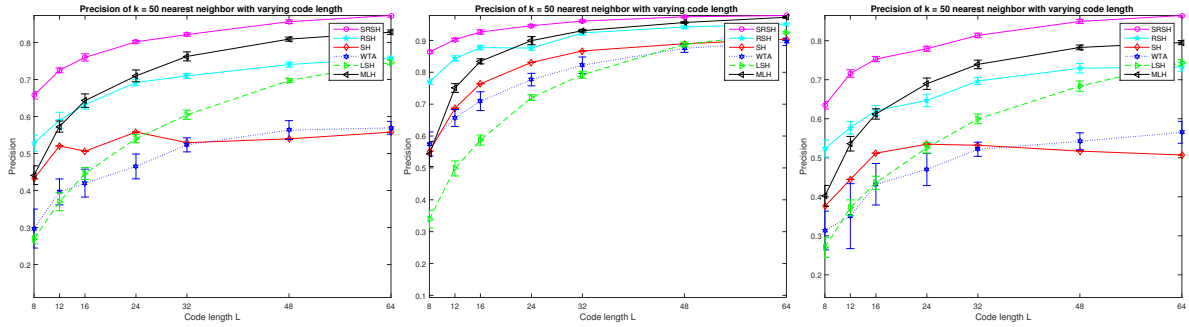


(a) LabelMe (b) MNIST (c) Peekaboom

*Figure 5.* Precision of 50 nearest neighbor retrieval based on Hamming distance.

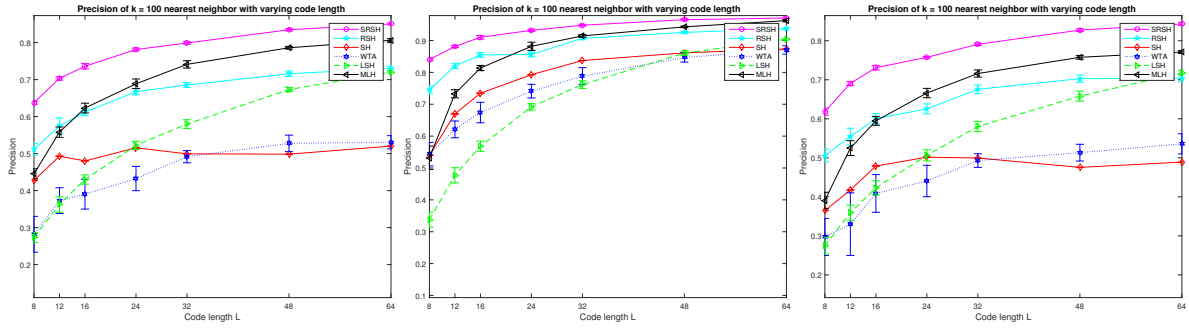


(a) LabelMe (b) MNIST (c) Peekaboom

*Figure 6.* Precision of 100 nearest neighbor retrieval based on Hamming distance.

*Vision, 2009 IEEE 12th International Conference on*, pp. 2130–2137. IEEE, 2009.

Li, Kai, Ye, Jun, and Hua, Kien A. What's making that sound? In *Proceedings of the ACM International Conference on Multimedia*, pp. 147–156. ACM, 2014.

Li, Xi, Lin, Guosheng, Shen, Chunhua, Hengel, Anton van den, and Dick, Anthony. Learning hash functions using column generation. *arXiv preprint arXiv:1303.0339*, 2013.

Liu, Wei, Wang, Jun, Kumar, Sanjiv, and Chang, Shih-Fu. Hashing with graphs. In *Proceedings of the 28th International Conference on Machine Learning (ICML-11)*, pp. 1–8, 2011.

Liu, Wei, Wang, Jun, Ji, Rongrong, Jiang, Yu-Gang, and Chang, Shih-Fu. Supervised hashing with kernels. In *Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conference on*, pp. 2074–2081. IEEE, 2012.

Liu, Wei, Mu, Cun, Kumar, Sanjiv, and Chang, Shih-Fu. Discrete graph hashing. In Ghahramani, Z., Welling, M., Cortes, C., Lawrence, N.D., and Weinberger, K.Q. (eds.), *Advances in Neural Information Processing Systems 27*, pp. 3419–3427. Curran Associates, Inc., 2014.

McAllester, David, Hazan, Tamir, and Keshet, Joseph. Direct loss minimization for structured prediction. In *Advances in Neural Information Processing Systems*, 2010.

Norouzi, Mohammad and Fleet, David J. Minimal loss hashing for compact binary codes. In *Proceedings of the 28th International Conference on Machine Learning (ICML-11)*, pp. 353–360, 2011.

Norouzi, Mohammad, Fleet, David J, and Salakhutdinov, Ruslan. Hamming distance metric learning. In *Advances in Neural Information Processing Systems*, pp. 1061–1069, 2012.

Ozuysal, Mustafa, Fua, Pascal, and Lepetit, Vincent. Fast keypoint recognition in ten lines of code. In *Computer Vision and Pattern Recognition, 2007. CVPR'07. IEEE Conference on*, pp. 1–8. Ieee, 2007.

Pele, Ofir and Werman, Michael. Robust real-time pattern matching using bayesian sequential hypothesis testing. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 30(8):1427–1443, 2008.

Salakhutdinov, Ruslan and Hinton, Geoffrey. Semantic hashing. *RBM*, 500(3):500, 2007.

Shakhnarovich, Gregory, Viola, Paul, and Darrell, Trevor. Fast pose estimation with parameter-sensitive hashing. In *Computer Vision, 2003. Proceedings. Ninth IEEE International Conference on*, pp. 750–757. IEEE, 2003.

Torralba, Antonio, Fergus, Robert, and Weiss, Yair. Small codes and large image databases for recognition. In *Computer Vision and Pattern Recognition, 2008. CVPR 2008. IEEE Conference on*, pp. 1–8. IEEE, 2008.

Tschopp, Dominique and Diggavi, Suhas. Approximate nearest neighbor search through comparisons. *arXiv preprint arXiv:0909.2194*, 2009.

Wang, Jingdong, Shen, Heng Tao, Song, Jingkuan, and Ji, Jianqiu. Hashing for similarity search: A survey. *CoRR*, abs/1408.2927, 2014.

Wang, Jun, Kumar, Sanjiv, and Chang, Shih-Fu. Semi-supervised hashing for large-scale search. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 34 (12):2393–2406, 2012.

Weiss, Yair, Torralba, Antonio, and Fergus, Rob. Spectral hashing. In *Advances in neural information processing systems*, pp. 1753–1760, 2009.

Yagnik, Jay, Strelow, Dennis, Ross, David A, and Lin, Ruei-sung. The power of comparative reasoning. In *Computer Vision (ICCV), 2011 IEEE International Conference on*, pp. 2431–2438. IEEE, 2011.